Multiagent Collaboration in Directed Improvisation

Barbara Hayes-Roth, Lee Brownston, and Robert van Gent

Knowledge Systems Laboratory
Stanford University
701 Welch Road, Bldg C
Palo Alto, CA 94304

Telephone: 415 723-0506
email: hayes-roth@cs.stanford.edu
fax: 415 725-5850

Abstract

Directed improvisation is a new paradigm for multiagent interaction. One or more human users direct one or more computer characters with scripted or interactive directions. The characters work together to improvise a course of behavior that follows the directions, expresses their distinctive individual styles, honors social conventions, and meets other objectives. The resulting "performance" reflects the collaboration among all of the human and computer agents. Directed improvisation has several attractive properties as a paradigm for multiagent human-computer interaction, which we illustrate in our testbed application, an animated improvisational theater company for children. Directed improvisation also presents distinctive requirements for agent interaction (emphasizing: situated, spontaneous, opportunistic behavior; very intimate interaction with shared control; and process-oriented evaluation criteria), which make it a useful addition to the domain of inquiry for multiagent systems.

Directed Improvisation

Directed improvisation is the simultaneous invention and performance of a new "work" under the constraints of user-specified directions. Besides its role in theater [7, 18, 22], directed improvisation mediates other activities: jazz music [9], planning and control of everyday behavior [12, 14, 16, 17], reactive behavior [1, 16], conversation [21], human-machine communication [26], scientific investigation [24], children's planning and playcrafting [2, 3, 23], and life management [6]. Indeed, most human behavior and interaction appears to incorporate some degree of directed improvisation.

We are studying directed-improvisation as a paradigm for multiagent human-computer interaction. Here, the "new work" is a course of behavior enacted by computer characters. Users give characters abstract directions, either interactively or in preconceived scripts. The characters improvise a collaborative course of behavior that follows the directions, expresses their distinctive individual styles, reflects social principles, and achieves other performance objectives. Thus, the characters obey their users, but also surprise and engage them along the way with artfully improvised behavior.

Section 2 of this paper illustrates directed improvisation in our current testbed application, an animated improvisational theater company for children. Section 3 presents an agent architecture to support improvisational characters. Section 4 discusses distinctive requirements for multiagent interaction that arise in the directed-improvisation paradigm.

Illustrative Application

We are developing an animated improvisational theater company to support children's participation in a variety of creative, playful, and educational activities [13, 17, 19]. Here we focus on characters and directed improvisation.

Each character in the company has a repertoire of physical and verbal behaviors (Figure 1), which serve as building blocks for improvisation. Different instances within a class allow characters to express life-like variability and moods. For example, a character may go to a destination by beeline or hop, depending on whether it feels determined or playful. Idiosyncratic instances allow characters to express life-like individual differences. For example, character C2 uses more casual verbal behaviors than C1. Even with nominally equivalent physical behaviors, characters may move differently due to differences in size and other physical features.

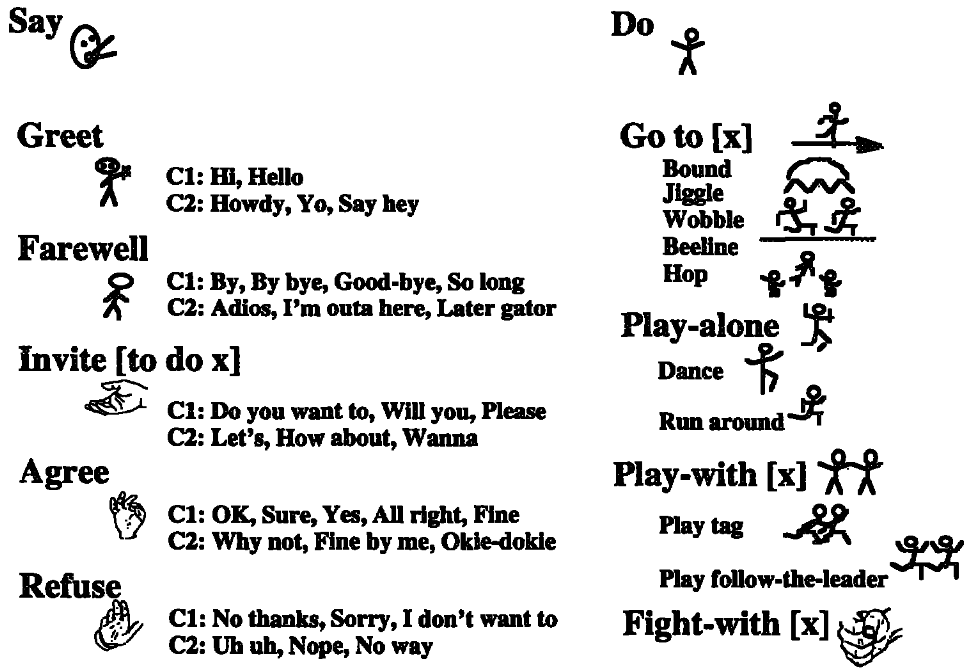

Figure 1. Excerpts from behavioral repertoires of two characters, C1 and C2.

Our illustrations are based on two characters that we adapted from the "woggles" system of J. Bates of Carnegie-Mellon University (including the RAL system of Production Systems Technologies). Unlike the CMU woggles, our characters have multiple gaits, speak lines, follow directions, and improvise. Each character has about 10 classes of physical behaviors and 20 classes of verbal behaviors, each with 1-5 instances. The characters' verbal behaviors and personalities were conceived and recorded by Aaron and Nora Hayes-Roth, who were 13 and 10 years old at the time. We plan to develop new human-like animated characters that will have similar behaviors, as well as new behaviors. To create stories, children direct the characters to adopt moods and to perform sequences of behaviors. The characters improvise within the constraints of the directions. We have conceived several interaction modes [13, 18], but present only two here for illustration.

In *animated-puppets mode*, children direct characters interactively by making choices from *situated behavior menus*. For example, in Figure 2, the large character feels cheerful and energetic. It considers going somewhere, playing alone, or inviting its partner to play. The small character feels OK, but tired. It considers going to the rest area, playing alone, or speaking. One child directs the large character to invite its partner to play. The character improvises an appropriate course of behavior (e.g., approach the small character, greet it, and say "Do you want to play follow the leader?"). Its executed behaviors change the shared situation and, as a result, the behaviors both characters subsequently consider. The children work side by side, directing their characters' moods and behavior to create a shared story, just as they would with physical puppets. The characters also collaborate on the playcrafting, improvising within the constraints of the children's directions. We have implemented animated puppets, with graphical user interfaces similar to the schematic ones in Figure 2, for the two characters described above [19].

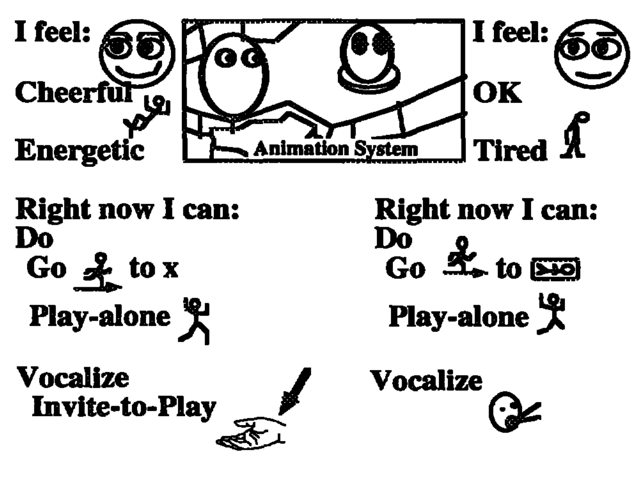

Figure 2. Snapshot from situated behavior menus for a two-character play.

In *animated actors mode*, children direct characters prospectively by giving them *synchronized behavior scripts*. For example, in Figure 3, the characters begin by playing independently (e.g., dancing, hopping on the pedestal). C2 determines when "awhile" has passed and suggests a game (e.g., hide and seek, follow the leader). C1 agrees and the two characters play the game together. The characters also may improvise script-independent behaviors (e.g., exchange greetings in a chance encounter while playing alone). We have implemented a few simple behavior scripts for brief conversations (e.g., greeting and reply; invitations and acceptances to play) and for games (e.g., follow the leader). We are developing a repertoire of such scripts as a form of reusable knowledge to be incorporated in different characters.

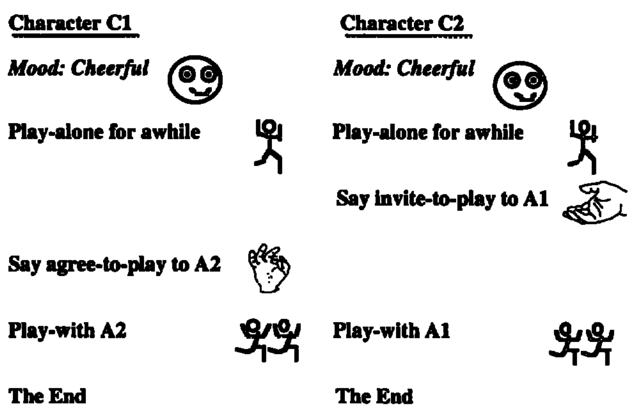

Figure 3. Synchronized Behavior Script.

The animated theater company illustrates multiagent human-computer collaboration in the directed-improvisation paradigm. Children collaborate to create the narrative structures of stories and direct the behavior of characters within the structure. Characters collaborate with children to implement the directions and with one another to determine the details of their interactions. Characters' improvisations may reflect moods and personalities, social conventions, relationships amd recent interactions, and improvisational expertise. The balance of creative work shifts among children and characters, depending on the specificity of the children's directions.

An Architecture for Improvisational Actors

Each animated character embodies an intelligent agent (Figure 4), whose two-level architecture comprises a *cognitive controller* and a *physical controller* [11, 17]. An agent's cognitive controller receives perceptual inputs from its physical controller, constructs an evolving model of its situation and its interactions with other characters, plans its physical and verbal behaviors, and sends control plans to its physical controller. An agent's physical controller interprets and filters inputs from its user interface and from sensors on its animated embodiment and sends those perceptions to its cognitive controller. It receives control plans from its cognitive controller and enacts those plans with appropriate physical behaviors, sending their outputs to effectors on its animated embodiment and to its user interface.

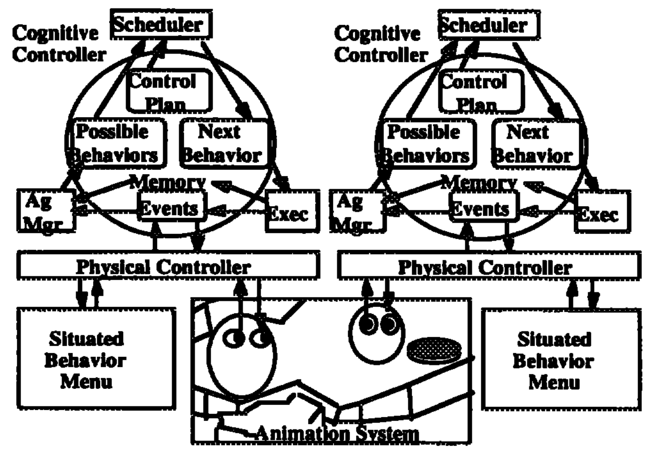

Figure 4. Framework for a Two-Character Production.

A *dynamic control model* [10, 12, 15, 16] underlies an agent's cognitive and physical controllers, each of which iterates this cycle: (a) its agenda manager triggers behaviors for the current situation; (b) its scheduler chooses the best triggered behavior based on the current control plan; and (c) its executor performs the chosen behavior. The cycle has three key features, as illustrated in Figure 5. First, the agent continually notices possible behaviors. For example, the agent might notice that it's able to play any of several games, go somewhere, say something, etc. Second, the agent describes its intended behavior in abstract control plans that permit alternative realizations. For example, the control plan "go to x" can be realized by "bound to pedestal," "hop to pedestal," "hop to rest area," etc. Third, the agent generates and modifies its control plans at run time. For example, the agent refines its plan to "go to x" by planning to "hop to pedestal." At each point in time, the agent performs possible behaviors that match its current control plans.

Our architecture is a natural framework for directed improvisation. The agenda manager generates situated behaviors—the raw materials out of which a character creates meaningful behavior patterns. Dynamic control plans can represent children's scripted and interactive directions. Dynamic control planning enables characters to improvise the details of directions and to augment them with undirected behaviors on the fly. Characters can use knowledge of the distinctive properties of behavior instances within a class to choose among them. For example, given the direction "Go to x," a character might choose any of five gaits.

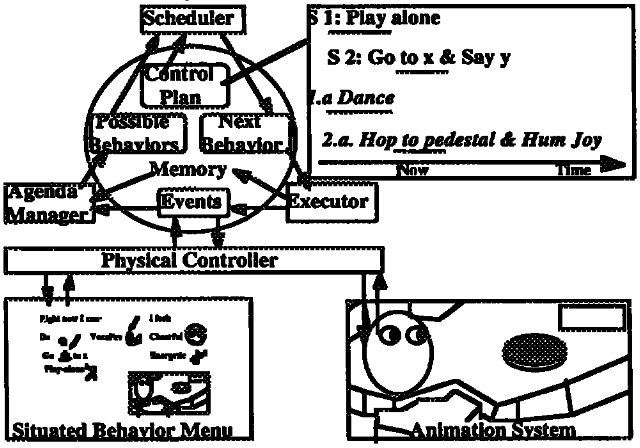

Figure 5. Dynamic Control Plans Integrate Users' Scripted and Interactive Directions with Characters' Improvisational Decisions (in *italics*).

If it previously had adopted a "silly" mood, it would choose a silly gait: hop, wobble, or jiggle. If the direction were more specific, "Go quickly to x," the character would choose a fast gait: bound or beeline. These examples illustrate a very simple form of directed improvisation, solo choice among alternative behaviors.

Table 1. Forms of Directed Improvisation.

Solo Improvisation	*Collaborative Improvisation*
One-Step Improvisation	
Choose among alternative logically equivalent behaviors Direction: Go to pedestal Improvise: Hop to pedestal-3	Respond to a partner's behaviors Partner: Greeting Improvise: Return greeting
Sequential Improvisation	
Construct a coherent path to a dramatic moment Direction: Play alone, Rest Improvise: Play alone, Get tired, Rest	Recognize and coordinate with a partner's behavior sequence Partner: Going toward pedestal Improvise: Go toward pedestal, Meet at pedestal
Patterned Improvisation	
Instantiate an improvisational schema Direction: Dance: Improvise: Iterate(Hop, twirl)	Recognize and participate in a partner's schema Partner: Play hide and seek? Improvise: I count to 10, etc.

Table 1 characterizes six forms of directed improvisation that vary in: (a) the involvement of partners; and (b) the complexity of the recognized situation and the improvised course of behavior. We are working on developing these more sophisticated forms of improvisational expertise in our agents.

To support multi-agent collaborative improvisation, we are elaborating our agent architecture so that agents can monitor, interpret, and predict their partners' behavior with the same mechanisms they use to represent, plan, monitor, and control their own behavior. In Figure 6, the large character enters the set first and makes a plan to look for something (1). The small character enters (2), observes its partner standing still, and infers nothing about its mood or plans (3). The small character plans to "start something" by playing alone for awhile and then hiding (4). Meanwhile, the large character observes its partner enter and start playing alone and infers that it is acting "shy" (5). The large character plans an appropriate interaction: approach, greet, and invite the shy partner to play (6). Now, the two characters have unknowingly constructed conflicting plans. One of them must change its plan and perhaps its model of its partner. However, as discussed below, it doesn't matter which character changes, as long as their observable interactions are plausible. Thus, if the large character succeeds in inviting the small character to play, the small character must abandon its plan and agree to play. If the small character notices the large character approaching and rushes over to tell it to hide, the large character must abandon its plan and its model of its partner and go hide. The extended architecture enables characters to generate and adapt their models of one another and their plans for their own behavior as their interaction unfolds.

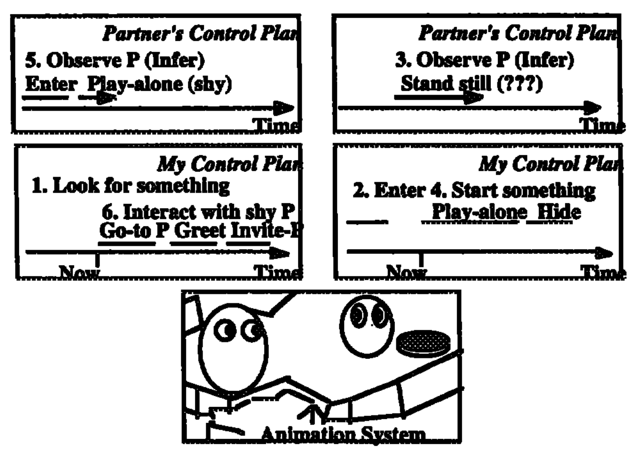

Figure 6. Simple Illustration of Collaborative Improvisation

Issues in Multiagent Interaction

Directed improvisation is a new interaction paradigm for multiagent systems. Although it shares important requirements and approaches with previously-studied multiagent paradigms (e.g., cooperative problem solving, distributed work, discretionary cooperation [8]), it presents three distinctive requirements that make it a useful addition to the domain of inquiry. We have taken these requirements from the writings and teachings of improvisation experts [7, 20, 22, 24] and view them as design objectives in our efforts to elaborate and instantiate the proposed agent architecture.

Directed improvisation demands behavior that is situated, spontaneous, and opportunistic in the service of abstract and weakly-constrained goals.

Most multi-agent paradigms assume that agents' individual and collective activities are predominantly and specifically goal-directed. In cooperative problem solving, the team's goal is to solve a shared problem. In distributed work, the team's goal is to get all of the work done as efficiently and effectively as possible. In discretionary cooperation, individual agents have their own goals and cooperate to the degree that doing so enhances achievement of their own goals or at least does not compromise them. Techniques for these paradigms focus on goal-directed reasoning in support of individual and group planning. Unplanned behavior occurs only when it is necessary to compensate for plan failure or when it happens to advance planned efforts to achieve established goals.

The directed-improvisation paradigm also assumes that agents have goals, but they are less well defined and less constraining than in other paradigms. Agents collaborating in a production share a general goal to produce a successful joint performance that meets the constraints of the users' directions. Although directions can vary in specificity, effective application of the paradigm typically involves relatively abstract directions that only weakly constrain the performance and give characters plenty of freedom to improvise. The working assumption is that producing the directed behaviors is easy; the art lies in the improvisation. The literature on improvisation reflects this assumption in its prescription of two underlying cognitive heuristics for the improviser: (a) *welcome possibilities* (Just let the words flow. Do not fear mistakes. Turn off the censor. Look for relationships. Do not plan too far ahead.); and (b) *pursue promising possibilities* (Make the natural response. Relate present actions to past actions. Keep the action on stage. Listen and respond to your partner. Take it to the extreme. Accept (don't block) offers). In contrast to the forward-looking, goal-directed reasoning and planning of agents in traditional paradigms, effective improvisers engage in backward-looking efforts to *reincorporate* incidental themes and behavioral qualities that they or their partners happen to have generated previously. Thus, the improviser at work is sometimes described as "a man walking backwards, trying to make sense of where he's been" [24]. In sum, the individual agent's behavior is firmly situated in the dynamic context, spontaneous in its short-term etiology, and opportunistic in its thematic relationships to other aspects of the performance. Ensemble behavior builds incrementally out of individual agents' actions and reactions. Achievement of the "goal" is not the specific product of a deliberate, provably correct process, but an emergent and uncertain epiphenomenon of the agents' real-time interactions.

Directed improvisation demands intimate collaboration and shared control among agents engaged in closely intertwined and interdependent behaviors.

Most multi-agent paradigms assume that agents have limited interest in working together and limited interactions during the actual performance of their collective activities. In cooperative problem solving and distributed work, agents work together because they are committed to solving a shared problem or to completing a shared job. The prevailing interaction model is to divide-and-conquer the joint task so that individual agents can work more or less independently. Multi-agent planning is used to allocate and coordinate tasks among agents, with run-time communication used primarily to exchange selected results of individual activities and, in some cases, to reallocate responsibilities. In discretionary cooperation, individual agents have their own goals and may or may not be willing to work together at all, depending on what it costs them. Competition for resources or conflict among goals may impede cooperation or even motivate agents to thwart one another's efforts, raising issues of trust and deception. Communications focus mainly on determining agents' willingness to cooperate. In all of these paradigms, research concentrates on the key question of how agents should decide in advance who is able and willing to do what. Techniques have been developed for: problem decomposition, communication of assumptions, beliefs, decisions, and commitments; negotiation and persuasion; conflict resolution and consensus building; organization of effective chains of command; and establishing mutually beneficial social laws.

By contrast, the directed improvisation paradigm assumes that agents are 100% committed to collaboration on their joint performance and that they will work together every step of the way. Their prospective work is so intertwined (and unpredictable) that agents do not even think of dividing it up ahead of time. Instead, effective improvisers rely on one another to do what is necessary to generate the work jointly and interactively in real time. The work is intrinsically collaborative: none of the participating agents can progress without interacting with its partners. Individual agents may introduce plot devices or instantiate improvisational routines. They may build small individual plans of activity in particular situations. However, all such structures, routines, and plans are tentative and dispensable under the fundamental rule of collaborative improvisation: *Accept all offers*. An "offer" is any explicit or implicit assertion. No matter how an individual agent feels about a partner's offer, no matter where the individual had been planning to go in the performance prior to the offer, no matter how the partner's offer redirects the individual's behavior, there can be only one response: "Yes." An effective improviser always embraces a partner's offers and tries to advance them with constructive, collaborative behavior. Since all of the participants are allowed to make offers and since not all offers are detected or understood by partners, each improviser must be willing to both lead and follow, with no preconception of when, how often, or how long. The dynamism and mutual adaptation of good improvisers reflects this underlying willingness to share control. Unlike control regimes based on organization, negotiation, turn-taking, or other explicit arrangements, *shared control* is an implicit arrangement in which the participants readily and immediately adopt and contribute to one another's assertions, goals, strategies, and tactics so that they can move forward together.

Directed improvisation succeeds when it produces a joint performance that follows the script and directions in an engaging manner.

Most multi-agent paradigms are "product-oriented"—they evaluate the objective consequences of agents' behavior against high objective standards: achievement of agents' individual and collective goals, optimal use of resources, and acceptability of side effects. These criteria for evaluating group performance imply related criteria for evaluating individual performance. To succeed as a member of the group, an agent must: reason correctly in performing its own tasks, reliably execute its planned physical behaviors, model its partners' knowledge and behavior correctly, make rational decisions about commitments to cooperative behavior, keep its commitments or at least inform its partners of changes, etc. Features of agents' behavior that do not affect the product are not valued, may be distracting, and, in the worst case, may carry unacceptable costs.

By contrast, the directed improvisation paradigm is "process-oriented"—the joint course of behavior enacted by computer characters is their product. Other than meeting the constraints of users' directions (which is usually assumed to be easy), there is no "correct" or "incorrect" performance. There is no concern with resources and unexpected side effects are simply folded into the performance. In general, there can be many alternative, equally "successful" performances of a given script—that is, performances that follow the directions in an engaging manner. In fact, in domains like our computer theater, children (and adults) may find it especially charming to observe or participate in repeated performances of a favorite story or one they have created themselves just to see how the characters will improvise anew. By implication, instead of behavior that is correct, rational, and reliable, effective improvisers produce behavior that appears appropriate in context, varies in different performances, and, in the best case, is endearingly idiosyncratic. The individual qualities that agents bring to a production are not costly distractions, but powerful sources of texture and depth in their contributions to the joint performance. In contrast to the all-business mentality of the ideal agent in traditional multi-agent paradigms, effective improvisers bring believable characters to life.

Concluding Remarks

In this paper, we have tried to illustrate the promise of directed improvisation as a paradigm for multiagent systems. First, it provides a structured framework for multiagent, human-computer collaboration [8]. Second, it explicitly allows run-time flexibility in the manner in which objectives can be achieved [1, 11, 16, 17]. Third, it offers a familiar, life-like, and potentially delightful interactive experience [4, 5]. Finally, because of its distinctive requirements for agent interaction, directed improvisation represents a useful addition to the domain of inquiry for multiagent systems.

References

1. Agre, P.E., and Chapman, D. What are plans for? Robotics and Autonomous Systems, 6, 17-34, 1990.

2. Applebee, A.N. The User's Concept of Story. Chicago: University of Chicago Press, 1978.

3. Baker-Sennett, J. Matusov, E., and Rogoff, B. Sociocultural processes of creative planning in children's playcrafting. In P. Light and G. Butterworth (eds.), Context and Cognition: Ways of Learning and Knowing. New York: Harvester Wheatsheaf, 1992.

4. Bates, J. The role of emotion in believable agents. Communications of the ACM, 1994, 37, 122-125.

5. Bates, J., Hayes-Roth, B., Laurel, B., and Nilsson, N. Working notes of the AAAI Symposium on believable agents, Menlo Park, CA: AAAI, 1994.

6. Bateson, M.C. Composing a Life. NY: The Atlantic Monthly Press, 1989.

7. Belt, L., and Stockley, R. Improvisation through Theater Sports. Seattle: Thespis Productions, 1991.

8. Charib-draa, B., Mandiau, R., and Millot, P. Distributed artificial intelligence: An annotated bibliography. SigArt Bulletin, 3,1992, 20-37.

9. Dean, R. Creative Improvisation: Jazz, Contemporary Music and Beyond. Philadelphia: Open University Press, 1989.

10. Hayes-Roth, B. A blackboard architecture for control. Artificial Intelligence, 1985, 26, 251-321.

11. Hayes-Roth, B. An architecture for adaptive intelligent systems. Artificial Intelligence, in press, 1994.

12. Hayes-Roth, B. Architectural foundations for real-time performance in intelligent agents. Real-Time Systems, 2, 99-125, 1990.

13. Hayes-Roth, B. Directed improvisation: A new paradigm for computer games. Proc. of the Computer Game Developers' Conference, Santa Clara, 1995.

14. Hayes-Roth, B. Opportunistic control of action. IEEE Trans. on Systems, Man, and Cybernetics, 12, 1575-1587 1993.

15. Hayes-Roth, B., and Collinot, A. A satisficing cycle for real-time reasoning in intelligent agents. Expert Systems with Applications, 7, 1993.

16. Hayes-Roth, B., and Hayes-Roth, F. A cognitive model of planning. Cognitive Science, 1979, 3, 275-310.

17. Hayes-Roth, B., Pfleger, K., Lalanda, P., Morignot, P., and Balabanovic, M. A domain-specific software architecture for adaptive intelligent systems. IEEE Trans. on Software Engineering, in press, 1994.

18. Hayes-Roth, B. Sincoff, E., Brownston, L., Huard, R., and Lent, B. Directed improvisation. Stanford University: Knowledge Systems Laboratory Technical Report KSL-94-61, 1994.

19. Hayes-Roth, B. Sincoff, E., Brownston, L., Huard, R., and Lent, B. Directed improvisation with animated puppets. Proc. of CHI '95 Conference on Computer-Human Interaction; Denver, 1995.

20. Johnstone, K. IMPRO: Improvisation in the Theatre. New York: Penguin Books, 1987.

21. Moore, J. Invited presentation at the National Conference on Artificial Intelligence, 1994.

22. Nachmanovitch, S. Free Play: Improvisation in life and art. Los Angeles: Jeremy P. Tarcher, Inc. 1990.

23. Rogoff, B., Gauvain, M., and Gardner, W. The development of children's skills in adjusting plans to circumstances. In S. Friedman, E. Scholnick, and R. Cocking (eds.), Blueprints for Thinking: The role of planning in psychological development. New York: Cambridge University Press, 1987, 303-20.

24. Ryan, P. Personal communication, Stanford University, Course HS 173 Improvisation - Discovering Spontaneity, 1994.

25. Stefik, M. Planning and meta-planning (MOLGEN: part 2). Artificial Intelligence, 16, 1981, 141-169.

26. Suchman, L. A. Plans and Situated Actions: The Problem of Machine/Human Communication. Cambridge, England: Cambridge University Press, 1987.